arXiv:1407.8154v2 [physics.comp-ph] 28 May 2015

A splitting approach for the Kadomtsev–Petviashvili equation[☆]

Lukas Einkemmer[a,*], Alexander Ostermann[a]

[a] *Department of Mathematics, University of Innsbruck, Austria*

Abstract

We consider a splitting approach for the Kadomtsev–Petviashvili equation with periodic boundary conditions and show that the necessary interpolation procedure can be efficiently implemented. The error made by this numerical scheme is compared to exponential integrators which have been shown in Klein and Roidot (SIAM J. Sci. Comput., 2011) to perform best for stiff solutions of the Kadomtsev–Petviashvili equation. Since many classic high order splitting methods do not perform well, we propose a stable extrapolation method in order to construct an efficient numerical scheme of order four. In addition, the conservation properties and the possibility of order reduction for certain initial values for the numerical schemes under consideration is investigated.

1. Introduction

The Kadomtsev–Petviashvili equation (KP equation) is a model of nonlinear wave propagation which was proposed in [11]; it is usually stated in the following form

$$\left(u_t + 6uu_x + \varepsilon^2 u_{xxx}\right)_x + \lambda u_{yy} = 0, \tag{1}$$

where λ and ε are two parameters that are determined by the physical problem under consideration. The KP equation appears in the description of long wavelength waves, where we choose either $\lambda = 1$ (weak surface tension) or $\lambda = -1$ (strong surface tension). In accordance with the literature (see, for example, [12]) we call the latter the KP I model and the former the KP II model.

The KP equation is a nonlinear dispersive partial differential equation that can be considered as a two-dimensional generalization of the well known Korteweg–de Vries equation (KdV equation). Similar to the KdV equation the KP equation is Hamiltonian and as a consequence does not include any dissipation. It exhibits many interesting physical phenomena such as soliton solutions and blow-up in finite time (see, for example, [13]).

[☆] This work is supported by the Austrian Science Fund (FWF) – project id: P25346.

[*] Corresponding author

Email address: lukas.einkemmer@uibk.ac.at (Lukas Einkemmer)

Before a numerical scheme is applied equation (1) is usually rewritten in evolution form

$$u_t + 6uu_x + \varepsilon^2 u_{xxx} + \lambda \partial_x^{-1} u_{yy} = 0, \tag{2}$$

where ∂_x^{-1} is to be understood as the regularized Fourier multiplier of $-\mathrm{i}/k_x$. That is, as in [12], we impose periodic boundary conditions and use the Fourier multiplier

$$\frac{-\mathrm{i}}{k_x + \mathrm{i}\lambda\delta},$$

where δ is equal to machine epsilon (the smallest number that in the finite precision arithmetic system under consideration yields a result different from one when added to one). That is, for the double precision floating point numbers employed in the simulations presented here, we have $\delta = 2^{-52}$.

The KP equation shows a number of interesting phenomena including soliton solutions and the appearance of small scale oscillations. For soliton solutions the stiffness of the KP equation is usually only a minor concern (since the linear part of the equation can be solved by spectral methods). In this setting various types of IMEX methods are usually very efficient. However, a number of phenomena do display stiff behavior and therefore pose a significant challenge for numerical schemes. In [12] it was found that in this context exponential integrators, in many instances, outperform IMEX and implicit Runge–Kutta methods. Furthermore, explicit time integrators suffer from a severe stability restriction of the time step size (due to the third derivative that appears in the dispersive term) that renders them computationally unfeasible.

In this paper we will demonstrate that splitting methods provide a viable and computationally attractive alternative to exponential integrators for stiff solutions of the KP equation. In section 2 we introduce the Strang splitting approach and an exponential integrator of order two. In section 3 the performance of the Strang splitting scheme is compared to that of the exponential integrator. The conservation properties of the splitting approach are investigated in section 4. In section 5 we consider both traditional high-order splitting schemes as well as a computationally attractive alternative approach based on Richardson extrapolation. Let us duly note that this approach avoids the stability problems often present if local Richardson extrapolation is applied to a nonlinear problem. We then provide, in section 6, a comparison of the run time between the second and fourth order methods for a given accuracy. In section 7 we consider order reduction that is observed for initial values which violate a constraint. In fact, we observe that the commonly observed order reduction is not present for the Strang splitting scheme. Finally, we conclude in section 8.

2. Numerical approach

In this paper we will exclusively employ the setting described in [12]. That is, the KP equation for a given initial value and periodic boundary conditions is propagated in time. Within this framework we are limited to initial values that are either periodic or decrease sufficiently fast for large values of $|x|$ and $|y|$.

In this setting, the form of equation (2) suggests an approach where the linear part can be solved very efficiently by means of fast Fourier techniques. This eliminates the (severe) stability constraint imposed by both the third and second order differential operators present in the KP equation. Exponential integrators (see, for example, [9]) exploit the fact that the linear part can be efficiently diagonalized. Similar to Runge–Kutta methods, time integrators of arbitrary order can be constructed where it is only required that the Burgers' nonlinearity can be evaluated efficiently. However, while this method manages to overcome a number of difficulties inherent in the numerical integration of the KP equation, it also suffers from a number of disadvantages due to the fact that the discretization of the Burgers' nonlinearity is essentially explicit. In fact, there is no mathematical proof that shows that exponential integrators are stable for the KP equation. Note that such results have been established for unbounded nonlinearities (using the parabolic smoothing property) and for bounded nonlinearities (see, for example, [9]). From a numerical standpoint, such considerations are important if the nonlinear dynamics is equally important as the linear dynamics (for example, if u is large in magnitude). In addition, it is often not clear how exponential integrators behave with respect to the conservation of invariants of the continuous system. For many interesting problems splitting methods solve both of these problems. In fact, stability (and convergence) of Strang splitting for a number of dispersive equations with a Burgers' nonlinearity are available in [10].

In the splitting approach considered here, we compute an approximate solution to the partial flows given by

$$u_t = Au = -\varepsilon^2 u_{xxx} - \lambda \partial_x^{-1} u_{yy} \tag{3}$$

and

$$u_t = B(u) = -6uu_x. \tag{4}$$

If it is possible to efficiently compute sufficiently accurate approximations to the partial flows given by equation (3) and (4), respectively, splitting methods constitute a viable approach. For example, the Strang splitting scheme for the step size τ is given by

$$u_{n+1} = \mathrm{e}^{\frac{\tau}{2}A} \left(\varphi_\tau^B \left(\mathrm{e}^{\frac{\tau}{2}A} u_n \right) \right),$$

where (for a given initial value u_n) the linear partial flow corresponding to equation (3) is denoted by $\mathrm{e}^{\tau A} u_n$ and the nonlinear partial flow corresponding to equation (4) is denoted by $\varphi_\tau^B(u_n)$. Let us also note that since the linear half-steps can be combined, the Strang splitting scheme does only need to compute the action of each partial flow once during each time step.

Before proceeding, let us note that the partial flow given in (4) is not well defined for arbitrarily large time steps. This is due to the fact that Burgers' equation develops a singularity for finite times. In principle this implies a step size restriction for the splitting approach. In all the simulations conducted in this paper the step size is determined by accuracy considerations only. However,

this would be a more serious concern in the small dispersion limit (i.e. where $\varepsilon \to 0$). We will not consider this case here but remark that due to the large gradients in the solution, a different numerical procedure for both time and space discretization seems to be in order then.

In the next section we will compare the splitting approach outlined above to the exponential integrator of order two given by

$$u_{n+1} = \mathrm{e}^{\tau A} u_n + \tau \varphi_1(\tau A) B(u_n) + \tau \varphi_2(\tau A) \left(B(U) - B(u_n) \right), \tag{5}$$

where

$$U = \mathrm{e}^{\tau A} u_n + \tau \varphi_1(\tau A) B(u_n)$$

and the φ_i functions are given by the recurrence relation

$$z\varphi_{k+1}(z) = \varphi_k(z) - \varphi_k(0)$$

with initial value $\varphi_0(z) = \mathrm{e}^z$.

3. Performance considerations

In [12] it has been argued that the splitting approach is not viable as the interpolation necessary to solve Burgers' equation (i.e., to compute an approximation to the action of φ_τ^B) is too costly compared to the computation of uu_x which only requires two Fast Fourier Transforms (FFTs) as well as some complex arithmetics. Therefore, we will consider this point in more detail in this section.

The algorithm of Cooley and Tukey requires approximately $5n \log n$ floating point operations. Libraries, such as the Fastest Fourier Transform in the West (FFTW [6]) used in our implementation, provide very efficient implementations of the FFT (including optimizations using SSE and AVX[1] and the use of more advanced algorithms). On the other hand, using the method of characteristics, we can derive an expression for the exact solution of equation (4)

$$u_1(x) = u\left(0, x - 6\tau u_1(x)\right), \tag{6}$$

where $u_1(x)$ is the solution of (4) at time τ with initial value $u(0, \cdot)$. Note that for the KP equation y is a parameter in the above equation (that is, we have to compute an approximation to $u_1(x)$ for each grid point in the y-direction). The representation given here is still implicit in u_1 and can be solved by conducting a fixed-point iteration. In a practical numerical scheme, this fixed-point iteration has to be truncated after a finite number of iterations (henceforth denoted by i). Of course, the value of i has a substantial impact on the performance. The other ingredient necessary is an interpolation algorithm. Such an algorithm is

[1] The Streaming Single instruction, multiple data Extension (SSE) and the Advanced Vector extensions (AVX) are a collection of CPU instructions that can be utilized to accelerate code segments that applies the same operation to multiple floating point numbers.

required as we have to determine the value of $u(0, x_i - 6\tau u_1(x_i))$, for each grid point x_i. Let us further note that using the FFT algorithm is not possible in this case as the translation does depend on x_i itself; this fact implies that the resulting points are no longer equidistant. However, similar to semi-Lagrangian methods (see, for example, [19]) we can use a (local) polynomial or a spline interpolation of sufficiently high degree.

Let us now consider the efficiency of constructing and evaluating a spline approximation. Construction of a cubic spline requires $\mathcal{O}(n)$ (real) floating point operations (the cost of the tridiagonal matrix solver). The resulting polynomial is accurate of order four. To evaluate a polynomial then requires $4n$ floating point operations (where we count one addition and one multiplication as one operation). Thus, one would conclude that even for medium sized problems the floating point operations count favors the spline interpolation. However, once we consider an implementation in C++ the performance of this scheme is somewhat disappointing. For example, using the GNU scientific library[2] (GSL [3]) we need approximately 200 ms to construct the spline and 150 ms for each fixed-point iteration.

On the other hand performing two FFTs (as is required to compute the Burgers' nonlinearity) requires only 120 ms. The second order exponential integrator, in total, requires the evaluation of two nonlinearities and an additional 6 FFTs for the computation of the matrix functions, yielding a total cost of approximately 600 ms per time step, whereas the Strang splitting algorithm requires approximately $320 + i \cdot 150$ ms per time step. Thus the Strang splitting scheme has approximately equal cost if we choose $i = 2$ (a value that is presumably too small). We have also used the ALGLIB library and found its performance significantly worse than GSL. Note that the GSL library also offers an interpolation that constructs and evaluates the interpolation polynomial for a number of specified grid points. We have used this to perform interpolation with a stencil of four grid points that is centered at the evaluation point. However, using this polynomial approximation approach from GSL does not significantly improve performance either; even though in this case we do not have to construct a global spline.

The issue here is not only one of optimization (GSL is most certainly not as well tuned as FFTW is) but in fact does relate to the problem that is being solved. The FFT algorithm must assume that it operates on an equidistant grid. This is not true for a spline or polynomial interpolation. In fact, all libraries require both an array of the grid points and the function values. Also, the FFT expansion is global thus alleviating the need for (a possible expensive) modulo operation in order to determine which part of the approximation needs to be accessed. Furthermore, GSL does not know a priori that we only employ fixed degree polynomials. Thus, it has to implement an algorithm that is stable even if high degree polynomials need to be constructed on a highly irregular grid. None of the difficulties listed above (except for performing a modulo operation)

[2] on an Intel Core i5-3427 CPU and a problem of size $2^{11} \cdot 2^9$.

are relevant here. Based on the Lagrange form, we have implemented a cubic approximation (i.e., an approximation of order four) that only requires $60 \cdot i$ ms. Note that for this implementation the cost of the cubic interpolation is equal to the cost of performing a single FFT. This then means that $i = 8$ would yield a Strang splitting scheme that is equal in execution time to the exponential integrator of order two. The details of this implementation are given in Appendix A.

Now, at least two questions remain to be answered: what value of i is required in order to obtain a sufficiently accurate approximation and how does the error of the Strang splitting scheme compare to the second order exponential integrator. To that end, we have conducted numerical experiments for the KP I and KP II equations using the Schwartzian initial value (as is done in [12], for example) given by

$$u(0, x, y) = -\tfrac{1}{2}\partial_x \mathrm{sech}^2\left(\sqrt{x^2 + y^2}\right). \tag{7}$$

The results are shown in Figure 1 (KP I equation) and Figure 2 (KP II equation). We observe that the Strang splitting scheme yields an error that is smaller by a factor of 10 for the KP I equation and smaller by a factor of 3 for the KP II equation. These gains can be realized by only performing three fixed-point iterations. The increase in accuracy together with the very competitive run-time leads to the conclusion that splitting methods can in fact be very competitive in the setting considered.

The results of the relative performance between the Strang splitting method and the exponential integrator of order two can be understood by considering the relative strength of the dispersive term εu_{xxx} and the Burgers' nonlinearity $6uu_x$. For the Schwartzian initial value the Burgers' nonlinearity is larger in magnitude by approximately a factor of 10. As time evolves dispersive effects eventually take over. This happens more slowly in the case of the KP I equation than for the KP II equation, which in turn explains the larger gain in accuracy achieved by the splitting approach in the former case (it is expected that splitting methods provide increased relative performance, compared to exponential integrators, as the importance of the Burgers' nonlinearity increases).

Let us further note that, as stated in [12], the analysis conducted above is strictly speaking only correct if the Fourier multipliers can be precomputed. This holds true for a constant step size integrator but not if adaptive step size control is employed. In fact, recomputing the Fourier multipliers (due to the complex exponential) is by at least a factor of 5 more costly than performing the forward and backward FFT. In the Strang splitting scheme this only affects a single exponential while in the exponential integrator of order two the Fourier multiplier for two additional φ functions have to be recomputed.

To conclude this section let us note that in performing the splitting algorithm spectral convergence is lost. This is due to the fact that we employ a polynomial interpolation in solving the Burgers' equation which is only of order four. For a fixed time step size the error in space is shown in Figure 3.

Let us note that while spectral convergence in space is certainly a desirable

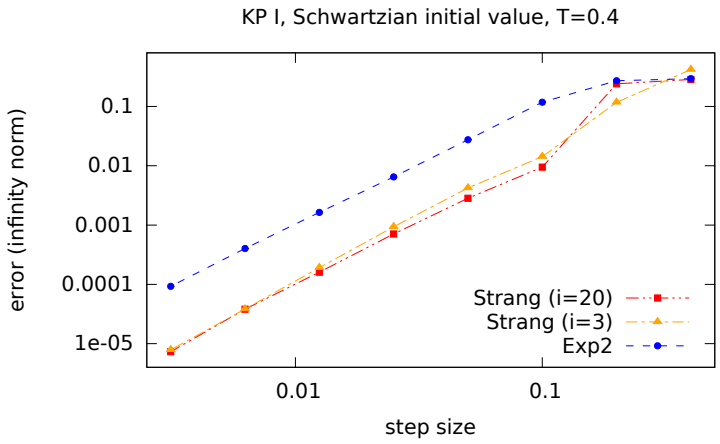


Figure 1: The error (in the infinity norm) as a function of the step size is shown at time $t = 0.4$ for the KP I equation using the Schwartzian initial value (7). The parameter ε is chosen equal to 0.1. To discretize space we have employed 2^{11} grid points in the x-direction and 2^9 grid points in the y-direction (on a domain of size $[-5\pi, 5\pi] \times [-5\pi, 5\pi]$). The number of iterations conducted to solve Burgers' equation for the Strang splitting scheme is denoted by i and the exponential integrator (5) of order two is referred to as Exp2. The error is computed using a reference solution with step size equal to 10^{-3}.

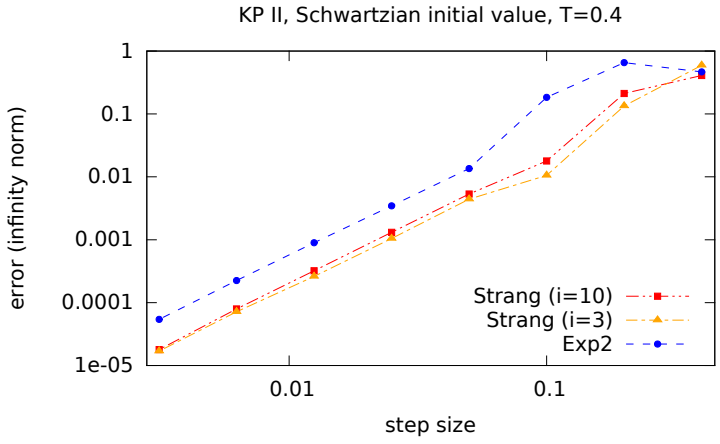


Figure 2: The error (in the infinity norm) as a function of the step size is shown at time $t = 0.4$ for the KP II equation using the Schwartzian initial value (7). The parameter ε is chosen equal to 0.1. To discretize space we have employed 2^{11} grid points in the x-direction and 2^9 grid points in the y-direction (on a domain of size $[-5\pi, 5\pi] \times [-5\pi, 5\pi]$). The number of iterations conducted to solve Burgers' equation for the Strang splitting scheme is denoted by i and the exponential integrator (5) of order two is referred to as Exp2. The error is computed using a reference solution with step size equal to 10^{-3}.

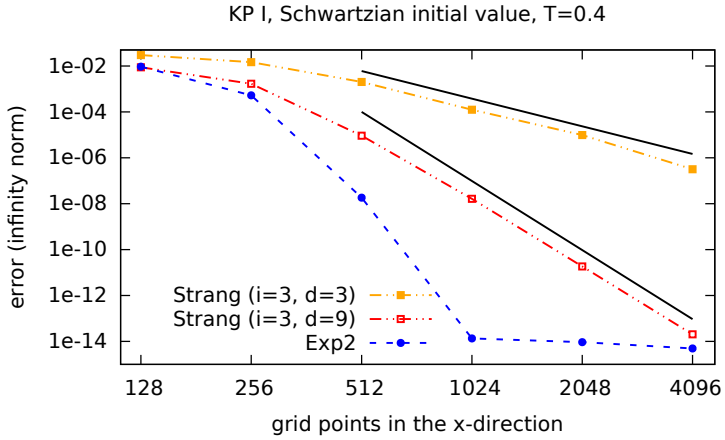


Figure 3: The error (in the infinity norm) as a function of the number of grid points in the x-direction is shown at time $t = 0.4$ for the KP I equation using the Schwartzian initial value (7). The parameter ε is chosen equal to 0.1. We employ 2^9 grid point in the y-direction and a fixed time step size that is equal to $\tau = 10^{-2}$. The domain under consideration is of size $[-5\pi, 5\pi] \times [-5\pi, 5\pi]$. The number of iterations conducted to solve Burgers' equation for the Strang splitting scheme is denoted by i and the degree of the polynomial interpolation is denoted by d. The exponential integrator (5) of order two is referred to as Exp2. The error is computed using a reference solution with 2^{13} grid points in the x-direction and a line of slope 4 and 10 is shown for comparison.

property, it has to be considered in the context of the time discretization error. Certainly there is no point in using a space discretization that is exact up to machine precision while making a time discretization error on the order of 10^{-2}. If a high accuracy in space is required polynomials of higher degree can be used. For example, the ninth degree polynomial interpolation shown in Figure 3 is approximately three times as costly as the cubic interpolation. Nevertheless, for $i = 3$ the Strang splitting scheme is still almost twice as fast compared to the exponential integrator of order two (for an equal number of grid points).

4. Conservation properties

In addition to using a scheme of sufficient accuracy at minimal computational cost, it is often desirable to employ a method that conserves certain invariants of the continuous problem (in this case the KP equation). This both ensures a physically consistent solution and usually facilitates the long time integration. It has long been known that an infinite number of quantities is conserved by the KP equation [16]. Note, however, that most of the high order invariants are only formal. That is, they are not well defined on suitable function spaces (see, e.g. [18]).

In this paper only linear and quadratic invariants are considered that have a clear physical interpretation. Following [17] these are the linear invariant $m(t)$

(corresponding to mass)

$$m(t) = \int_\Omega u(t,x,y)\,\mathrm{d}(x,y)$$

and the quadratic invariant $M(t)$ (corresponding to momentum)

$$M(t) = \int_\Omega u(t,x,y)^2\,\mathrm{d}(x,y).$$

In addition, for the KP equation the constraint

$$\int_{-\infty}^{\infty} \partial_{yy} u(t,x,y)\,\mathrm{d}x = 0 \tag{8}$$

is satisfied. This property, however, is respected for both the Strang splitting scheme as well as the exponential integrator up to machine precision (a consequence of the regularization).

Since Runge–Kutta methods preserve linear invariants (such as the mass in the KP equation) we might expect that the same holds true for exponential Runge–Kutta methods (all of the exponential integrators considered in this paper are in fact exponential Runge–Kutta methods). A more formal definition (see [9]) is given in Definition 1.

Definition 1. A exponential Runge–Kutta method applied to $u' = Au + B(t,u)$, $u(0) = u_0$ is given by

$$\begin{aligned} u_1 &= \mathrm{e}^{\tau A} u_0 + \tau \sum_{i=1}^{s} b_i(\tau A) G_i, \\ U_i &= \mathrm{e}^{c_i \tau A} u_0 + \tau \sum_{j=1}^{s} a_{ij}(\tau A) G_j, \\ G_j &= B(c_j \tau, U_j), \end{aligned}$$

where u_1 is an approximation to $u(\tau)$. The method is said to have $s \in \mathbb{N}$ stages and is uniquely determined by the coefficients $c_i \in \mathbb{R}$ and the coefficient functions b_i and a_{ij}, where $i, j \in \{1, \ldots, s\}$. The functions b_i are assumed to be linear combinations of φ_k functions.

In contrast to Runge–Kutta methods, we have to assume that the linear invariant under consideration is conserved for both the flow generated by A and the flow generated by B. This assumption is satisfied for the KP equation. In the following theorem we assume that A and B already have been discretized in space in such a way that the linear invariant considered is a conserved quantity of the discretized system. It should, however, be duly noted that the proof of Theorem 2 can just as well be carried out for the case where space is left continuous.

Theorem 2. *An exponential Runge–Kutta method preserves every linear invariant that is preserved by both the flow generated by A and the flow generated by B.*

Proof. Since the b_i are linear combination of φ_k functions (which, in general, will be evaluated for different step sizes), in order to show that d is an invariant of the numerical method we have to show that

$$d^{\mathrm{T}} u_1 = d^{\mathrm{T}} u_0.$$

For the exponential Runge–Kutta method we have

$$d^{\mathrm{T}} u_1 = d^{\mathrm{T}} u_0 + \tau \sum_{i=1}^{s} d^{\mathrm{T}} b_i(hA) G_i,$$

since $d^{\mathrm{T}} \mathrm{e}^{\tau A} u_0 = d^{\mathrm{T}} u_0$.

Now, we will show that $v(\tau) = \tau^k \varphi_k(\tau A) g$ satisfies

$$v'(\tau) = A v(\tau) + \frac{\tau^{k-1}}{k!} g, \qquad v(0) = 0. \tag{9}$$

Let us recall the recurrence relation for the φ_k function

$$\varphi_{k+1}(\tau A) g = (\tau A)^{-1} \left(\varphi_k(\tau A) - \varphi_k(0) \right) g$$

for which upon multiplication by τ^{k+1} and differentiating with respect to time we get

$$\begin{aligned}
\partial_\tau \tau^{k+1} \varphi_{k+1}(\tau A) g &= A^{-1} \partial_\tau \left(\tau^k \varphi_k(\tau A) - \tau^k \varphi_k(0) \right) g \\
&= \tau^k \varphi_k(\tau A) g + \tau^{k-1} A^{-1} \left(\frac{1}{(k-1)!} - k \varphi_k(0) \right) g \\
&= A (\tau A)^{-1} \left(\tau^{k+1} \varphi_k(\tau A) - \tau^{k+1} \varphi_k(0) \right) g + \frac{\tau^k}{k!} g \\
&= A \left(\tau^{k+1} \varphi_{k+1}(\tau A) g \right) + \frac{\tau^k}{k!} g.
\end{aligned}$$

A simple calculation in the case for φ_1 completes the induction.

Since we can assume that both $d^{\mathrm{T}} A w = 0$ and $d^{\mathrm{T}} g = 0$ for any w and g, we immediately follow from equation (9) that

$$d^{\mathrm{T}} v(t) = d^{\mathrm{T}} v(0) = 0$$

which implies that

$$d^{\mathrm{T}} \tau b_i(\tau A) G_i = 0.$$

This completes the proof. □

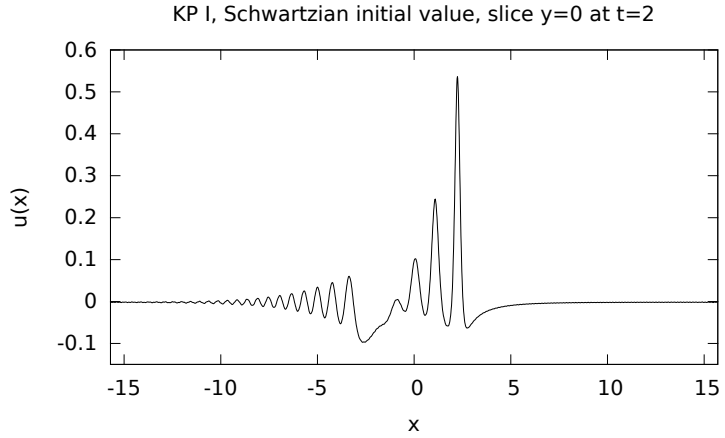


Figure 4: A one-dimensional slice (at $y = 0$) of the numerical solution of the KP I equation for the Schwartzian initial value (7) at $t = 2$ is shown. To discretize space 2^{11} grid points are employed in the x-direction and 2^9 grid points are employed in the y-direction (on a domain of size $[-5\pi, 5\pi] \times [-5\pi, 5\pi]$).

Before continuing let us note that all the methods considered in this paper or in [12] satisfy the assumption on the coefficient functions b_i given in Definition 1. Furthermore, since we employ a FFT based discretization in space, which conserves the mass exactly, we expect that the exponential integrator considered here do in fact conserve the mass (up to machine precision).

Now, let us numerically investigate the conservation of mass. To that end we perform simulations of the KP I and the KP II equation using the Schwartzian initial value up to the final time $t = 2$. A slice of the solution (for $y = 0$) is shown in Figure 4 (for the KP I equation) and in Figure 5 (for the KP II equation).

The error in the mass, that is $|m(t) - m(0)|$, is shown as a function of time in Figure 6 (for the KP I equation) and in Figure 7 (for the KP II equation). We observe, as expected from the theoretical result, that the exponential integrator conserves the mass up to machine precision, while for the Strang splitting scheme the error ranges from 10^{-6} to 10^{-10} depending on the number of iterations i performed and the number of grid points used. From the perspective of the splitting approach this behavior seems to be disappointing and perhaps contrary to intuition. However, it is entirely expected since by using the cubic polynomial interpolation we no longer have exponential convergence (as is the case for the Fourier approximation) and consequently a projection error is made in computing a solution to Burgers' equation. As we can see from Figure 6 the error does depend (weakly) on the number of fixed-point iterations conducted. Note that if we employ a finer space discretization (and increase the number of iterations) then the error in mass of the splitting approach does decrease as well (see Figures 6 and 7).

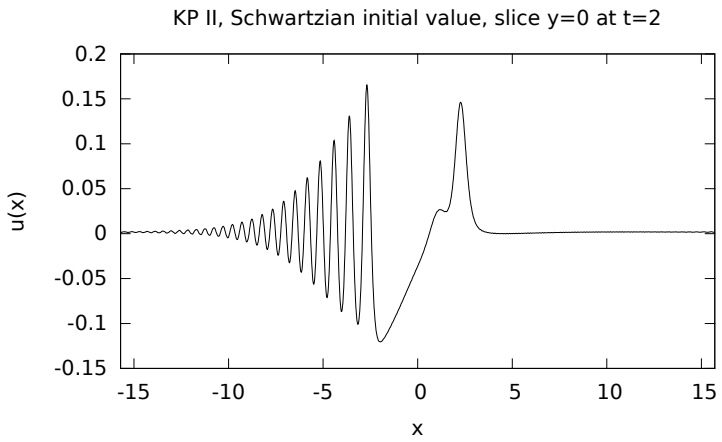


Figure 5: A one-dimensional slice (at $y = 0$) of the numerical solution of the KP II equation for the Schwartzian initial value (7) at $t = 2$ is shown. To discretize space 2^{11} grid points are employed in the x-direction and 2^9 grid points are employed in the y-direction (on a domain of size $[-5\pi, 5\pi] \times [-5\pi, 5\pi]$).

It is, however, not clear what the ramifications for long time integration are. In the context of semi-Lagrangian methods this was studied in some detail. It was found that even though the mass in such interpolation methods is not exactly conserved, they remain remarkably stable over long times (see, for example, [4]).

Now, let us consider the conservation of momentum. The results of the numerical solutions are shown in Figure 8 (for the KP I equation) and Figure 9 (for the KP II equation). In the former case we observe that for $2^{11} \times 2^9$ grid points the Strang splitting scheme is more accurate by two orders of magnitude, while in the latter case only a difference of one order of magnitude in accuracy is observed. Contrary to the exponential integrator, where the time step size has a significant impact on the conserved quantities, we observe a decrease in the error in mass as the number of grid points is increased. To obtain these results, we also have to slightly increase the number of iterations. The additional iterations performed, as compared to the order plots presented in the last section, do not appreciably decrease the error in mass but do result in better conservation properties (if a sufficiently fine space discretization is used).

5. High order splitting

An mathematical rigorous result (see, for example, [1]) shows that if real time steps are assumed and if the order of a splitting method is strictly larger than two, both partial flows have to be computed for a step size that is smaller than zero (i.e., we have to conduct steps backward in time). An alternative, see [8], is to use complex time steps (with positive real part). The former can

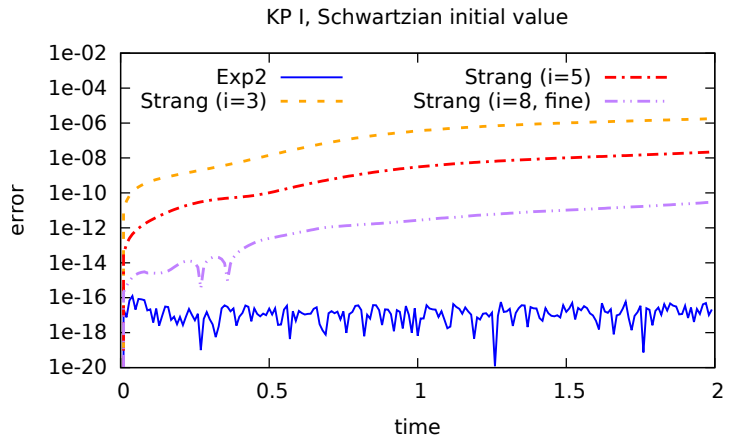


Figure 6: The error in mass, i.e. $|m(t) - m(0)|$ is shown as a function of time for the KP I equation. A time step of size 10^{-2} is used. To discretize space 2^{11} grid points are employed in the x-direction and 2^9 grid points are employed in the y-direction (on a domain of size $[-5\pi, 5\pi] \times [-5\pi, 5\pi]$), except for the fine discretization in which case $2^{13} \times 2^9$ grid points are used.

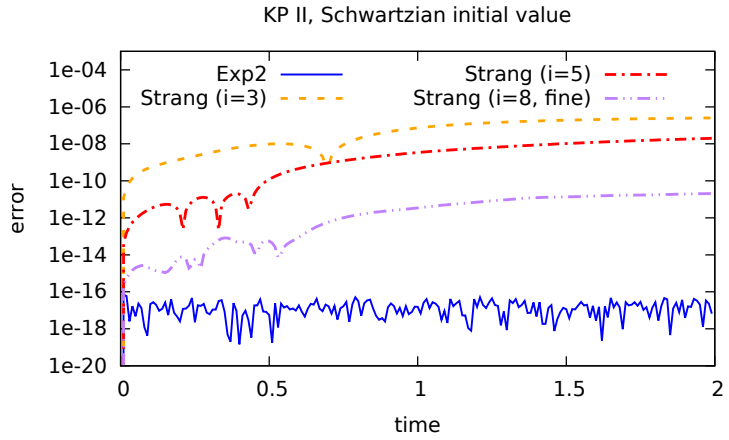


Figure 7: The error in mass, i.e. $|m(t) - m(0)|$ is shown as a function of time for the KP II equation. A time step of size 10^{-2} is used. To discretize space 2^{11} grid points are employed in the x-direction and 2^9 grid points are employed in the y-direction (on a domain of size $[-5\pi, 5\pi] \times [-5\pi, 5\pi]$), except for the fine discretization in which case $2^{13} \times 2^9$ grid points are used.

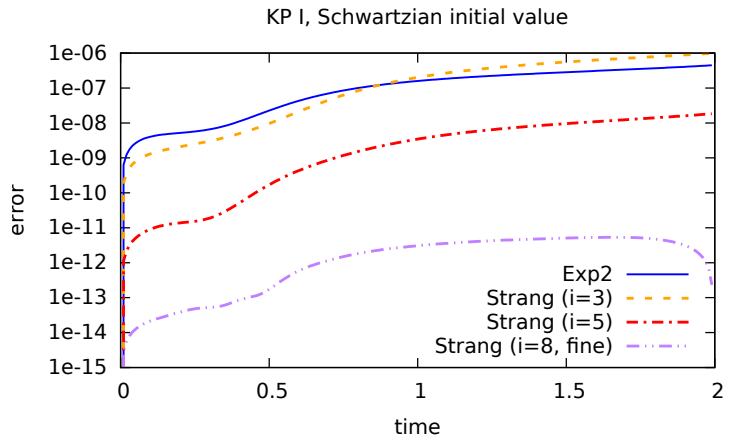


Figure 8: The error in momentum, i.e. $|M(t) - M(0)|$ is shown as a function of time for the KP I equation. A time step of size 10^{-2} is used. To discretize space 2^{11} grid points are employed in the x-direction and 2^9 grid points are employed in the y-direction (on a domain of size $[-5\pi, 5\pi] \times [-5\pi, 5\pi]$), except for the fine discretization in which case $2^{13} \times 2^9$ grid points are used.

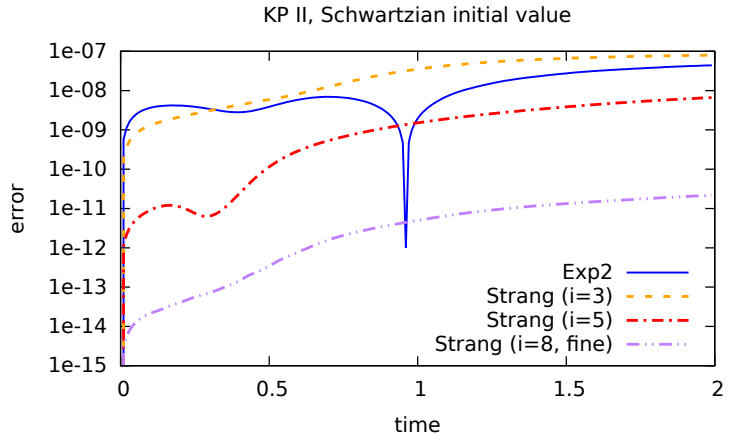


Figure 9: The error in the momentum, i.e. $|M(t) - M(0)|$ is shown as a function of time for the KP II equation. A time step of size 10^{-2} is used. To discretize space 2^{11} grid points are employed in the x-direction and 2^9 grid points are employed in the y-direction (on a domain of size $[-5\pi, 5\pi] \times [-5\pi, 5\pi]$), except for the fine discretization in which case $2^{13} \times 2^9$ grid points are used.

be used in purely hyperbolic partial differential equations to obtain methods of arbitrary order, while the latter can be used in purely parabolic partial differentials equations (with some performance penalty due to the necessity of using complex arithmetics) to obtain methods of high order.

The KP equation is hyperbolic and its eigenvalues are purely imaginary. Therefore, in principle, employing a splitting approach with negative time steps is a possibility. However, due to the regularization introduced in section 1 taking negative time steps results in an exponential amplification of round-off errors.

To explain this behavior let us consider the initial value problem

$$u_t(t,x,y) = \partial_x^{-1} u_{yy}(t,x,y)$$

which, after regularization, yields the following equations in Fourier space

$$\hat{u}_t(t,k_x,k_y) = \frac{-i}{k_x + i\lambda\delta} k_y^2 \hat{u} = \frac{-\lambda\delta - ik_x}{k_x^2 + (\lambda\delta)^2} k_y^2 \hat{u}(t,k_x,k_y).$$

Since these equations are decoupled in the wavevectors (k_x, k_y) we can consider $k_x = 0$ for an arbitrary wavevector $k_y \neq 0$. The solution is then given by

$$\hat{u}_t(t,0,k_y) = \mathrm{e}^{-tk_y^2/(\lambda\delta)} \hat{u}(0,0,k_y) \tag{10}$$

which for $t < 0$ implies that any non-zero value in any of these components is exponentially amplified which is clearly a undesirable behavior for any implementation in finite precision.

If the initial value does not satisfy the constraint given in equation (8) this regularization exponentially damps the modes with $k_x = 0$. This is in fact the expected behavior of the underlying continuous model. However, even if the constraint is satisfied by the initial value numerical errors (such as those made in the approximation of Burgers' equation) result in a non-zero value for these modes. Such a behavior, however, is unphysical and will eventually pollute the numerical solution. Thus, we need some mechanism to dissipate these modes if high order splitting schemes are to be applied.

Since the factor in the exponential of equation (10) immediately sets the modes under consideration to zero (even for very small time steps), we propose to implement this behavior implicitly in the numerical scheme. The corresponding solution is, up to machine precision, equivalent to the regularization procedure introduced in the introduction (which is used extensively in the literature). Note that the mode with $k_x = 0$ and $k_y = 0$ is constant in time and thus no regularization is required. At least for initial values that satisfy constraint (8) this allows us to perform negative time steps (the value of the modes under consideration is neither changed by a positive nor by a negative time step).

While the Strang splitting scheme is the universally employed second order splitting scheme, a variety of of different fourth order splitting schemes have been proposed. The often used triple jump scheme is relatively cheap from a computational point of view. Its implementation is only three times as expensive as an implementation of the Strang splitting scheme. Unfortunately, the triple

jump schemes employs negative time steps of length 1.7τ. This, implies that we require additional iterations in order to obtain good accuracy for the solution of Burgers' equation. However, even if this is done, numerical simulation suggest that the error constants of this method is very disappointing.

The length of the negative time step necessary for fourth order splitting methods can be reduced by considering additional stages. For example, the methods of Suzuki and McLachlan (see, for example, [7]) require roughly five times the computational effort compared to the Strang splitting scheme. In our numerical simulations the method of McLachlan has been found to perform best for both the KP I and KP II equation. However, on drawback of the method by McLachlan is that it is somewhat expensive from a computational point of view.

As an alternative, we consider the so-called Richardson extrapolation algorithm which enables the construction of higher order methods from an (almost) arbitrary numerical one-step method S_τ with step size τ. In the following we limit ourselves to the case where S_τ is a method of order two. We proceed by performing a step with length τ and two steps with $\tau/2$. Then, the final approximation u_{n+1} is computed from u_n as follows

$$u_{n+1} = \frac{4S_{\tau/2}\left(S_{\tau/2}(u_n)\right) - S_\tau(u_n)}{3}.$$

This procedure eliminates the leading error term in S_τ and due to the symmetry of the Strang splitting scheme results in a method which is consistent of order four. Note that since symmetry can be defined for the semi-discrete case (i.e. after space has already been discretized) order four is achieved independent of the space discretization under consideration. However, in general, the resulting scheme is not stable in the nonlinear case. In fact, we observe this lack of stability for the KP equation.

Therefore, we propose to apply a global extrapolation algorithm (see [20]). First, we compute

$$v_{n+1} = S_\tau(v_n)$$

and

$$w_{n+1} = S_{\tau/2}\left(S_{\tau/2}(w_n)\right),$$

where v_0 and w_0 are equal to the initial value u_0. Then, we compute the final approximation, for each time step n, as follows:

$$u_n = \frac{4w_n - v_n}{3}.$$

Note that this is in fact the Richardson extrapolation algorithm. But instead of applying it at each time step, we first compute a solution with time step τ and a solution with time step $\tau/2$ and then apply the extrapolation procedure independently for each time step. This alleviates the stability problems as both v_n and w_n are computed by the same unconditionally stable scheme (but using a different step size).

In [12] a number of exponential integrators have been compared in the context of the KP equation. It was found that the method of Cox and Matthews [2], the method of Krogstad [15], and the method of Hochbruck and Ostermann [9] do show almost identical performance characteristics (even though they differ in run-time as well as accuracy). We have chosen to compare the extrapolation scheme described in this section and the method by McLachlan with the method of Cox and Matthews. The results are shown in Figure 10 (for the KP I equation) and Figure 11 (for the KP II equation). We observe that in case of the KP I equation the method of McLachlan is more accurate by approximately an order of magnitude compared both to the extrapolation method and the method of Cox and Matthews. For the KP II equation the method of Cox and Matthews performs better and is is as accurate as the method of McLachlan. Both methods are superior to the extrapolation method by approximately a factor of 7. Note that the last statement is only true in the asymptotic case. There is a region (up to an error of approximately 10^{-2}) where the extrapolation method is more accurate.

The method of Cox and Matthews requires 4 evaluations of the nonlinearities and the computation of 12 matrix functions. This gives a total of 32 FFTs that have to be performed which compared to the second order method increases the cost by a factor of 3.2. This is almost the same increase in cost by a factor of 3 which is required for the extrapolation method. However, the method of McLachlan is 2.5 times as expensive as the extrapolation method. For a fourth order method this would require a gain in accuracy of almost a factor of 40 for the method to be competitive. Let us note, however, that the method of McLachlan shows better conservation properties compared to the extrapolation approach.

Furthermore, we have analyzed the conservation properties of the fourth order methods considered in this section in case of the KP I equation. As before we integrate the equation until final time $T = 2$. In this setting the method of McLachlan is the most robust scheme preserving momentum up to 10^{-10} for a time step size between 10^{-2} and $3 \cdot 10^{-2}$ (using the fine space discretization and $i = 8$). In this case lack of conservation is only due to the error made in solving Burgers' equation. On the other hand the extrapolation scheme shows a behavior similar to the method of Cox and Matthews. While for small time steps conservation to high accuracy can be observed (using a time step size of $\tau = 10^{-2}$ we observe an error in momentum approximately equal to $2 \cdot 10^{-12}$ for the extrapolation scheme and $6 \cdot 10^{-9}$ for the method of of Cox and Matthews), even if the time step size is only increased to $3 \cdot 10^{-2}$ the error in momentum increases to $2 \cdot 10^{-7}$ for the extrapolation scheme and to $3 \cdot 10^{-7}$ for the method of Cox and Matthews. The behavior of the extrapolation scheme is due to the fact that the error in momentum is now limited by the non-conservative nature of the extrapolation procedure.

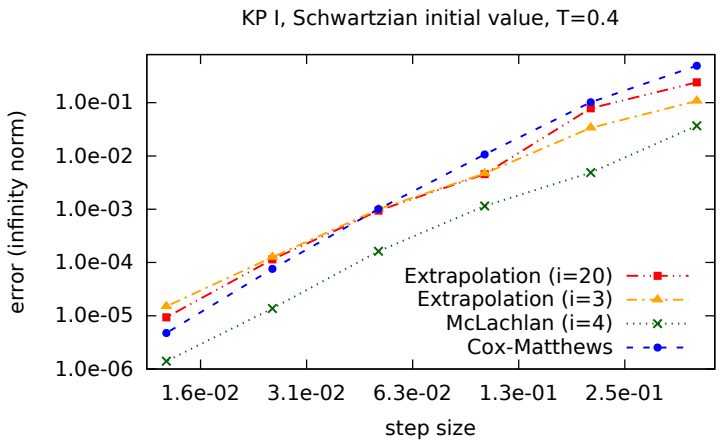


Figure 10: The error (in the infinity norm) as a function of the step size τ is shown at time $t = 0.4$ for the KP I equation using the Schwartzian initial value (7). The parameter ε is chosen equal to 0.1. To discretize space we have employed 2^{11} grid points in the x-direction and 2^9 grid points in the y-direction (on a domain of size $[-5\pi, 5\pi] \times [-5\pi, 5\pi]$). The number of iterations conducted to solve Burgers' equation for the extrapolation scheme and McLachlan's method is denoted by i. We have used the exponential integrator developed by Cox and Matthews. The error is computed using a reference solution with step size equal to 10^{-3}.

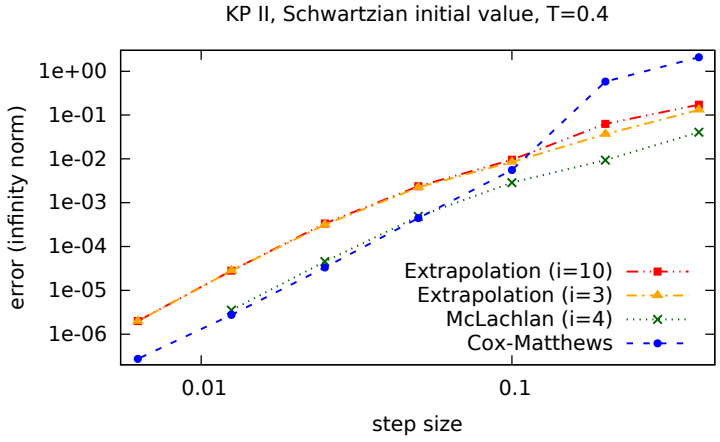


Figure 11: The error (in the infinity norm) as a function of the step size is shown at time $t = 0.4$ for the KP II equation using the Schwartzian initial value (7). The parameter ε is chosen equal to 0.1. To discretize space we have employed 2^{11} grid points in the x-direction and 2^9 grid points in the y-direction. The number of iterations conducted to solve Burgers' equation for the extrapolation scheme and McLachlan's method is denoted by i. We have used the exponential integrator developed by Cox and Matthews (on a domain of size $[-5\pi, 5\pi] \times [-5\pi, 5\pi]$). The error is computed using a reference solution with step size equal to 10^{-3}.

6. Performance comparison

In the previous sections we have only considered the error as a function of the time step size. Together with the performance considerations given in section 3, we are able to compare the relative performance of the second order exponential integrator and the Strang splitting scheme. However, it is difficult to compare a second to a fourth order scheme. In addition, it is instructive to compare the various schemes in terms of the run time that is necessary to achieve a given accuracy in time.

The purpose of this section is to perform the corresponding comparison. We employ the same initial values that are considered in section 3. The numerical results consider a tolerance between 10^{-1} and 10^{-5} and are shown in Figure 12 (for the KP I equation) and Figure 13 (for the KP II equation). For convenience the speedup of using the Strang splitting approach (which is superior to the extrapolation scheme for the accuracy considered here) to the best exponential integrator (either the second order exponential integrator or the method of Cox and Matthews) is indicated for a tolerance of 10^{-2} and 10^{-3}.

We observe that for low accuracy the Strang splitting scheme is faster by a factor of 3 to 7 compared to the exponential integrators. This is true for both the KP I and KP II equations. In addition, we observe that even for relatively low accuracy the method of Cox and Matthews is superior to the exponential integrator of order two considered here. On the other hand, the extrapolation method only overtakes the splitting method for a tolerance of approximately 10^{-5}. Let us further note that the performance of the extrapolation method is always better than that of the method by Cox and Matthews (although the difference between the two methods for accuracies below 10^{-5} is negligible). Note that even if we employ a ninth degree polynomial interpolation (instead of the cubic interpolation considered so far) the performance of the extrapolation method and the exponential integrator of Cox and Matthews is almost equal (the extrapolation method is faster for the KP I equation and the method of Cox and Matthews is slightly faster for the KP II equation).

7. Initial values that violate a constraint

It is well known that the KP equation does satisfy the constraint (8) for positive times $t > 0$ even if this is not the case for the initial value (see, for example, [5] and [18]). This behavior is enforced by a discontinuity in time for the continuous problem and by the regularization for the discrete problem. For time integration schemes this usually results in order reduction (see, for example, [14]).

To investigate this phenomenon we will consider the initial value

$$u(0, x, y) = \alpha \mathrm{e}^{-D(x^2+y^2)} \tag{11}$$

which does not satisfy the constraint and

$$u(0, x, y) = \beta x \mathrm{e}^{-D(x^2+y^2)} \tag{12}$$

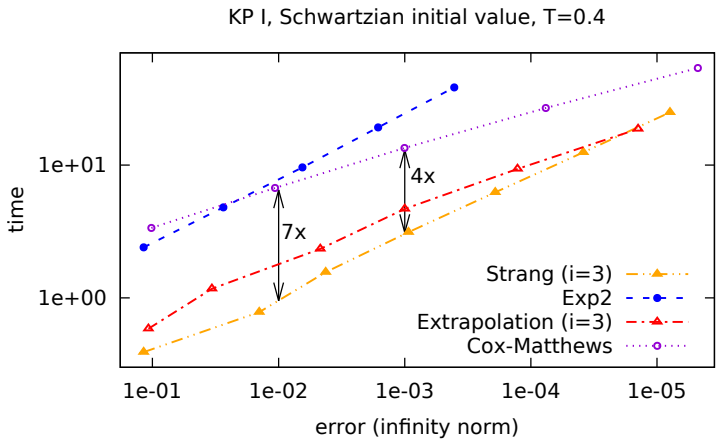


Figure 12: The error (in the infinity norm) as a function of the run time is shown at time $t = 0.4$ for the KP I equation using the Schwartzian initial value (7). The parameter ε is chosen equal to 0.1. To discretize space we have employed 2^{11} grid points in the x-direction and 2^9 grid points in the y-direction (on a domain of size $[-5\pi, 5\pi] \times [-5\pi, 5\pi]$). The number of iterations conducted to solve Burgers' equation for the Strang splitting scheme is denoted by i and the exponential integrator (5) of order two is referred to as Exp2. The error is computed using a reference solution with step size equal to 10^{-3}.

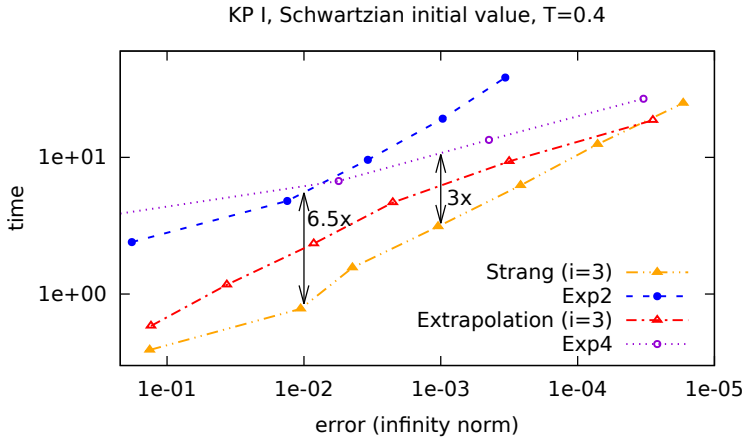


Figure 13: The error (in the infinity norm) as a function of the run time is shown at time $t = 0.4$ for the KP II equation using the Schwartzian initial value (7). The parameter ε is chosen equal to 0.1. To discretize space we have employed 2^{11} grid points in the x-direction and 2^9 grid points in the y-direction (on a domain of size $[-5\pi, 5\pi] \times [-5\pi, 5\pi]$). The number of iterations conducted to solve Burgers' equation for the Strang splitting scheme is denoted by i and the exponential integrator (5) of order two is referred to as Exp2. The error is computed using a reference solution with step size equal to 10^{-3}.

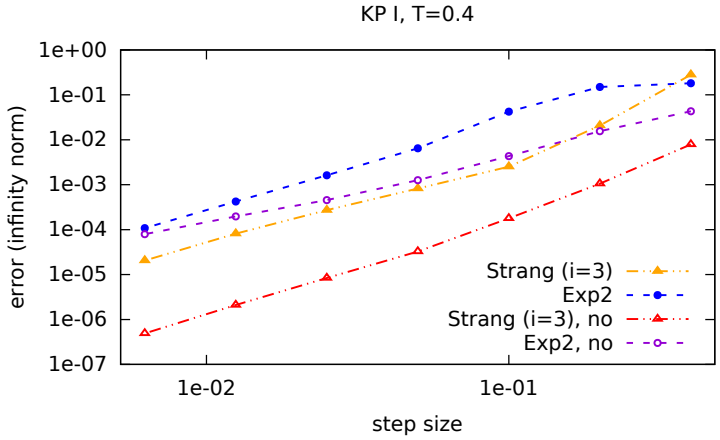


Figure 14: The error (in the infinity norm) as a function of the step size is shown at time $t = 0.4$ for the KP I equation using the initial values given in equation (11) and (12). The numerical results for the former initial value, which does not satisfy the constraint is denoted by an additional *no* in the legend of the plot. The parameter ε is chosen equal to 0.1. To discretize space we have employed 2^{11} grid points in the x-direction and 2^9 grid points in the y-direction (on a domain of size $[-5\pi, 5\pi] \times [-5\pi, 5\pi]$). The number of iterations conducted to solve Burgers' equation for the Strang splitting scheme is denoted by i and the exponential integrator of order two is referred to as Exp2. The error is computed using a reference solution with step size equal to 10^{-3}.

which satisfies the constraint. For both initial values $D = 0.5$ is used. The parameters $\alpha = 0.35$ and $\beta = 0.6$ have been chosen such that the maximal amplitude of both initial values is comparable.

The result of our numerical experiments are shown in Figure 14. For the second order exponential integrator we observe order reduction in case of the initial value that does *not* satisfy the constraint (even though this seems to be the easier problem). On the other hand, for the Strang splitting scheme the numerical results are consistent with a numerical method of order two (i.e. no severe order reduction is present). These results (as shown in Figure 14) imply that for the evaluation of the Gaussian pulse the Strang splitting scheme is more than two orders of magnitude more accurate compared to the exponential integrator of order two (for the same step size).

8. Conclusion & Outlook

We have demonstrated that splitting methods for the KP equation can be efficiently implemented and achieve performance that is significantly better compared to state of the art time integrators. The efficient implementation of the projections necessary for computing an approximation to the solution of Burgers' equation, as demonstrated in Appendix A, yields a Strang splitting scheme that, in addition to improved accuracy, is less expensive (as in most situations

we only have to compute a few fixed-point iterations) compared to the exponential integrator of order two. In addition, conservation of momentum is exact for the time integrator. A challenge for such methods is the error made by the space approximation which no longer shows spectral convergence in the number of grid points. Also conservation of mass and momentum is influenced by the polynomial approximation in space. We consider this as further research.

We have also considered high order splitting schemes (for a modified regularization) and proposed an extrapolation method for Strang splitting. The latter method is computationally attractive and achieves comparable accuracy to the exponential integrator of Cox and Matthews for the KP I equation but worse accuracy for the KP II equation. However, the good conservation of momentum observed for the Strang splitting scheme is lost by conducting the extrapolation. Note that the fourth order method by McLachlan shows the best conservation properties among the methods considered in this paper. It is, however, more expensive computationally.

Let us note that, for applications which require long time integration neither the exponential integrators (due to their conservation properties) nor the splitting approach (due to the lack of efficient high order methods) provide an ideal numerical scheme. We consider this as further research.

References

[1] S. Blanes and F. Casas. On the necessity of negative coefficients for operator splitting schemes of order higher than two. *Appl. Numer. Math.*, 54(1):23–37, 2005.

[2] S.M. Cox and P.C. Matthews. Exponential time differencing for stiff systems. *J. Comput. Phys.*, 176(2):430–455, 2002.

[3] M. Galassi et al. *GNU Scientific Library Reference Manual*. Network Theory Ltd, 3rd edition, 2009.

[4] F. Filbet and E. Sonnendrücker. Comparison of Eulerian Vlasov solvers. *Comput. Phys. Commun.*, 150(3):247–266, 2003.

[5] A.S. Fokas and L.Y. Sung. The Cauchy problem for the Kadomtsev–Petviashvili–I equation without the zero mass constraint. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 125, pages 113–138, 1999.

[6] M. Frigo. A fast Fourier transform compiler. In B.G. Ryder and B.G. Zorn, editors, *Proceedings of the 1999 ACM SIGPLAN Conference on Programming Language Design and Implementation, Atlanta, Georgia, USA, May 1-4, 1999*, pages 169–180. ACM, 1999. ISBN 1-58113-094-5.

[7] E. Hairer, C. Lubich, and G. Wanner. *Geometric numerical integration: structure-preserving algorithms for ordinary differential equations*. Springer-Verlag Berlin Heidelberg, 2nd edition, 2006.

[8] E. Hansen and A. Ostermann. High order splitting methods for analytic semigroups exist. *BIT Numer. Math.*, 49(3):527–542, 2009.

[9] M. Hochbruck and A. Ostermann. Exponential integrators. *Acta Numer.*, 19:209–286, 2010.

[10] H. Holden, C. Lubich, and N. Risebro. Operator splitting for partial differential equations with Burgers nonlinearity. *Math. Comp.*, 82(281):173–185, 2013.

[11] B.B. Kadomtsev and V.I. Petviashvili. On the stability of solitary waves in weakly dispersing media. *Sov. Phys. Dokl.*, 15:539–541, 1970.

[12] C. Klein and K. Roidot. Fourth order time-stepping for Kadomtsev-Petviashvili and Davey-Stewartson equations. *SIAM J. Sci. Comput.*, 33 (6):3333–3356, 2011.

[13] C. Klein and J.C. Saut. Numerical study of blow up and stability of solutions of generalized Kadomtsev–Petviashvili equations. *J. Nonlinear Sci.*, 22(5):763–811, 2012.

[14] C. Klein, C. Sparber, and P. Markowich. Numerical study of oscillatory regimes in the Kadomtsev–Petviashvili equation. *J. Nonlinear Sci.*, 17(5): 429–470, 2007.

[15] S. Krogstad. Generalized integrating factor methods for stiff PDEs. *J. Comput. Phys.*, 203(1):72–88, 2005.

[16] J. Lin and H.H. Chen. Constraints and conserved quantities of the Kadomtsev–Petviashvili equations. *Phys. Lett. A*, 89(4):163–167, 1982.

[17] A.A. Minzoni and N.F. Smyth. Evolution of lump solutions for the KP equation. *Wave Motion*, 24(3):291–305, 1996.

[18] L. Molinet, J.C. Saut, and N. Tzvetkov. Remarks on the mass constraint for KP type equations. *SIAM J. Math. Anal.*, 39(2):627–641, 2007.

[19] E. Sonnendrücker, R. Jean, P. Bertrand, and A. Ghizzo. The semi-Lagrangian method for the numerical resolution of the Vlasov equation. *J. Comput. Phys.*, 149(2):201–220, 1999.

[20] J.G. Verwer and H.B. de Vries. Global extrapolation of a first order splitting method. *SIAM J. Sci. Comput.*, 6(3):771–780, 1985.

Appendix A. Efficient implementation of cubic polynomial interpolation

The implementation shown here is based on the Lagrange form of the interpolation polynomial through four equidistant nodes. One important aspect of the algorithm is to determine the integer and fractional part of the evaluation

point with respect to the numerical grid. Usually this would involve modulo operations which, however, significantly impact performance. Therefore we only use casting to integer and replace modulo operations by arithmetic operations, wherever possible. The algorithm is divided into two loops; this approach yields a performance gain of 1.5 compared to the monolithic implementation (using the GCC C++ compiler). The code is shown in Algorithm 1.

Algorithm 1 Computation of the fixed-point iteration used to solve (6) (including the construction and evaluation of the Lagrange interpolation polynomial).

```
// Evaluates a cubic Lagrange polynomial on the grid {-1,0,1,2}.
double lagrange3(double x,double u0m1, double u0p0, double u0p1,
          double u0p2) {
  double lm1 = -0.16666666666666667*x*(x-1.0)*(x-2.0);
  double l0  =  0.5*(x+1.0)*(x-1.0)*(x-2.0);
  double l1  = -0.5*(x+1.0)*x*(x-2.0);
  double l2  =  0.16666666666666667*(x+1.0)*x*(x-1.0);
  return u0m1*lm1 + u0p0*l0 + u0p1*l1 + u0p2*l2;
}

// u0: the initial value, u1: the result of the computation
// nx/ny: number of grid points in x/y-direction, L: domain length
// fp_it: number of fixed-point iterations conducted
void burgers(array2d& u0, array2d& u1, double h, double L, int nx,
        int ny, int fp_it=3) {

  // Some values that can be precomputed.
  double adv = 6.0*h*double(nx)/L;
  vector<double> d_i(nx);
  for(int i=0; i<nx; i++)
    d_i[i] = double(i);
  vector<double> xred(nx);

  // Iteration over the y-direction.
  for(int j=0; j < ny; j++) {

    // The restrict keyword tells the compiler that there is no
    // pointer aliasing. This is essential for vectorization.
    double* __restrict _xred  = &xred[0];

    // The fixed-point iteration
    for (int k = 0; k < fp_it; k++) {

      // Avoid copying u0 to u1 in the first iteration.
      double* __restrict _us = (k==0) ? &u0(0,j) : &u1(0,j);

      // Determine an array of positions in [0,L].
      for(int i = 0; i < nx; i++) {
        double x = d_i[i] - adv*_us[i];
        if(x < 0)   x += nx;  // We assume that |u|<L/(6*h) holds.
        if(x >= nx) x -= nx;  // We assume that |u|<L/(6*h) holds.
        _xred[i] = x;
      }

      const double * __restrict _u0 = &u0(0,j);
            double * __restrict _u1 = &u1(0,j);
      for (int i = 0; i < nx; i++) {
        // Determine the interpolation nodes.
        int p0 = (int)_xred[i];
        if(p0 == nx)  p0 = 0;
        int pm1 = p0-1;
        if(pm1 == -1)  pm1 = nx-1;
        int p1 = p0+1;
        if(p1 == nx)  p1 = 0;
        int p2 = p1+1;
        if(p2 == nx)  p2 = 0;

        // Evaluate the Lagrange interpolation.
        double x = _xred[i] - double(p0);
        _u1[i] = lagrange3(x, _u0[pm1], _u0[p0], _u0[p1], _u0[p2]);
      }
    }
  }
}
```